arXiv:0708.2351v1 [cs.DS] 17 Aug 2007

Randomized algorithm for the k-server problem on decomposable spaces*

Judit Nagy-György†

February 1, 2008

Abstract

We study the randomized k-server problem on metric spaces consisting of widely separated subspaces. We give a method which extends existing algorithms to larger spaces with the growth rate of the competitive quotients being at most $O(\log k)$. This method yields $o(k)$-competitive algorithms solving the randomized k-server problem, for some special underlying metric spaces, e.g. HSTs of "small" height (but unbounded degree). HSTs are important tools for probabilistic approximation of metric spaces.

Keywords: k-server, on-line, randomized, metric spaces.

1 Introduction

In the theory of designing efficient virtual memory-management algorithms, the well studied paging problem plays a central role. Even the earliest operation systems contained some heuristics to minimize the amount of copying memory pages, which is an expensive operation. A generalization of the paging problem, called the k-server problem was introduced by Manasse, McGeoch and Sleator in [14], where the first important results were also achieved. The problem can be formulated as follows. Given a metric space with k mobile servers that occupy distinct points of the space and a sequence of requests (points), each of the requests has to be served, by moving a server from its current position to the requested point. The goal is to minimize the total cost, that is the sum of the distances covered by the servers; the optimal cost for a given sequence ϱ is denoted $\mathrm{opt}(\varrho)$.

An algorithm is online if it serves each request immediately when it arrives (without any prior knowledge about the future requests).

Definition 1. *An online algorithm A is c-competitive if for any initial configuration C_0 and request sequence ϱ it holds that*

$$\mathrm{cost}(A(\varrho)) \leq c \cdot \mathrm{opt}(\varrho) + I(C_0),$$

where I is a non-negative constant depending only on C_0.

---

[*] Research is partially supported by OTKA T049398.

[†] Department of Mathematics, University of Szeged, Aradi vértanúk tere 1, H-6720 Szeged, Hungary, email: Nagy-Gyorgy@math.u-szeged.hu

The competitive ratio of a given online algorithm A is the infimum of the values c with A being c-competitive. The k-server conjecture (see [14]) states that there exists an algorithm A that is k-competitive for any metric space. Manasse et al. proved that k is a lower bound [14], and Koutsoupias and Papadimitriou showed $2k-1$ is an upper bound for any metric space [12].

In the randomized online case (sometimes this model is called the oblivious adversary model [5]) the competitive ratio can be defined in terms of the expected value as follows:

Definition 2. *A randomized online algorithm R is c-competitive if for any initial configuration C_0 and request sequence ϱ we have*

$$\mathrm{E}(\mathrm{cost}(R(\varrho))) \leq c \cdot \mathrm{opt}(\varrho) + I(C_0),$$

where I is a non-negative constant depending only on C_0 and $\mathrm{E}(\mathrm{cost}(R(\varrho)))$ denotes the expected value of $\mathrm{cost}(R(\varrho))$.

The competitive ratio of the above randomized algorithm is defined analogously.

In the randomized version there are more problems that are still open. The randomized k-server conjecture states that there exists a randomized algorithm with a competitive ratio $\Theta(\log k)$ in any metric space. The best known lower bound is $\Omega(\log k/\log^2 \log k)$ [3]. A natural upper bound is the bound $2k+1$ given for the deterministic case. By restricting our attention to metric spaces with a special structure, better bounds can be achieved: Fiat et al. showed a lower bound $H_k = \sum_{i=0}^{k} i^{-1} \approx \log k$ for uniform metric spaces [10], which turned out to be also an upper bound, see McGeoch and Sleator, [13].

In this paper we also consider a restriction of the problem, namely we seek for an efficient randomized online algorithm for metric spaces that are "μ-HST spaces" [1, 2] and defined as follows:

Definition 3. *A μ-hierarchically well-separated tree (μ-HST) is a metric space defined on the leaves of a weighted, rooted tree T with the following properties:*

1. *The edge weight from any node to each of its children is the same.*
2. *The edge weights along any path from the root to a leaf are decreasing by the factor μ from one level to the next. The weight of an edge incident to a leaf is one.*

The μ-HST spaces play an important role in the so-called metric space approximation technique developed by Bartal [2]. Fakcharoenphol et al [11] proved that every weighted graph on n vertices can be α-probabilistically approximated by a set of μ-HSTs, for an arbitrary $\mu > 1$ where $\alpha = O(\mu \log n/\log \mu)$.

It has been shown in [15] that for any $2k$-HST with an underlying tree T that has a small depth and maximum degree there exists a polylog(k)-competitive randomized algorithm for the k-server problem. By slightly modifying the approach of Csaba and Lodha [7] and Bartal and Mendel [4] we show that there exists such an algorithm for any μ-HST that has a small depth and arbitrary maximum degree t, given $\mu \geq \min\{k, t\}$.

2 Notation

Suppose the points of a metric space can be partitioned into t blocks, $B_1, \ldots, B_t$, such that the diameter of each block is at most δ and whenever x and y are points of different blocks, their distance is exactly Δ. Suppose also that $\Delta/\delta = \mu \geq k$ holds. The above metric space is μ-decomposable [15].

For a given request sequence ϱ we denote its ith member by ϱ_i, and the prefix of ϱ of length i by $\varrho_{\leq i}$.

Given a block B_s, a request sequence ϱ having only requests from B_s and a number ℓ of servers inside B_s and an algorithm A, let $\mathrm{cost}(A_s(\ell, \varrho))$ denote the cost computed by the algorithm A for these inputs and $\mathrm{opt}_s(\ell, \varrho)$ (in the latter case also the initial position of the servers can be chosen). If ϱ is nonempty, $\mathrm{opt}_s(0, \varrho)$ is defined to be infinite.

Definition 4. *The* demand *of the block B_s for the request sequence ϱ that contains only requests from ϱ is*

$$D_s(\varrho) := \min\{\ell \,|\, \mathrm{opt}_s(\ell, \varrho) + \ell\Delta = \min_j\{\mathrm{opt}_s(j, \varrho) + j\Delta\}\},$$

if ϱ is nonempty, otherwise it is 0.

Visually, $D_s(\varrho)$ denotes the least number of servers to be moved into the initially empty block B_s to achieve the optimal cost for the sequence ϱ. We note that the behaviour of the sequence $D_s(\varrho_1), D_s(\varrho_{\leq 2}), \ldots, D_s(\varrho_{\leq i}), \ldots, D_s(\varrho)$ is unclear.

3 Algorithm X

In the rest of the paper we suppose that there exists a randomized online algorithm A and a function f with $\frac{f(\ell)}{\log \ell}$ being monotone increasing (constants are allowed), and we have

$$\mathrm{E}[\mathrm{cost}(A_s(\ell, \varrho))] \leq f(\ell) \cdot \mathrm{opt}_s(\ell, \varrho) + \frac{f(\ell) \cdot \ell\delta}{\log \ell} \tag{1}$$

for any ℓ and s. Having Algorithm A, we can define our "shell algorithm" X that uses A as a subroutine inside the blocks.

3.1 The Algorithm

The algorithm uses A as a subroutine and it works in phases. Let $\varrho^{(p)}$ denote the sequence of the pth phase. In this phase the algorithm works as follows:

> Initially we mark the blocks that contain no servers.
>
> When $\varrho_i^{(p)}$, the ith request of this phase arrives to block B_s, we compute the demand $D_s(\varrho_{\leq i}^{(p)})$ and the maximal demand
>
> $$D_s^*(\varrho_i^{(p)}) = \max\{D_s(\varrho_{\leq j}^{(p)}) | j \leq i\}$$
>
> for this block (note that these values do not change in the other blocks).
>
> – If $D_s^*(\varrho_i^{(p)})$ is less than the number of servers in B_s at that moment, then the request is served by Algorithm A, with respect to the block B_s.
>
> – If $D_s^*(\varrho_i^{(p)})$ becomes equal to the number of servers in B_s at that moment, then the request is served by Algorithm A, with respect to the block B_s and we mark the block B_s.
>
> – If $D_s^*(\varrho_i^{(p)})$ is greater than the number of servers in B_s at that moment, we mark the block B_s and perform the following steps until we have $D_s^*(\varrho_i^{(p)})$ servers in that block or we cannot execute the steps (this happens when all the blocks become marked):

- Let us choose an unmarked block $B_{s'}$ randomly uniformly, and a server from this block also randomly. We move this chosen server to the block B_s (such a move is called a *jump*), either to the requested point, or, if there is already a server occupying that point, to a randomly chosen unoccupied point of B_s. If the number of servers in $B_{s'}$ becomes $D^*_{s'}(\varrho^{(p)}_i)$ via this move, we mark that block. In both B_s and $B_{s'}$ we restart algorithm A from the current configuration of the block.

If we cannot raise the number of servers in block B_s to $D^*_s(\varrho^{(p)}_i)$ by repeating the above steps (all the blocks became marked), then Phase $p+1$ is starting and the last request is belonging to this new phase.

Our main result is the following:

Theorem 5. *Algorithm X is $c \cdot \log k \cdot f(k)$-competitive for some constant c.*

In the following two subsections we will give an upper bound for the cost of Algorithm X and several lower bounds for the optimal cost in an arbitrary phase. The above theorem easily follows from these.

For convenience we modify the request sequence ϱ in a way that does not increase the optimal cost and does not decrease the cost of any online algorithm, hence the bounds we get for this modified sequence will hold also in the general case. The modification is defined as follows: we extend the sequence by repeatedly requesting the points of the halting configuration of a (fixed) optimal solution. We do this till $\sum_{s=1}^{t} D^*_s(\varrho^{(u)}_{\leq i})$ becomes k. Observe that the optimal cost does not change via this transformation, and any online algorithm works the same way in the original part of the sequence (hence online), so the cost computed by any online algorithm is at least the original computed cost.

3.2 Upper bound

In the first step we prove an auxiliary result. We recall from [8] that an online matching problem is defined similarly to the online k-server problem with the following two differences:

1. Each of the servers can moved only once;

2. The number of the requests is at most k, the number of the servers.

For any phase p of Algorithm X we can associate the following matching problem MX. The underlying metric space of the matching problem is a finite uniform metric space that has the blocks B_s as points and a distance Δ between any two different points. Let $\hat{D}_s(p)$ denote the number of servers that are in the block B_s just at the end of phase p. Now in the associated matching problem we have $\hat{D}_s(p-1)$ servers originally occupying the point B_s. During phase p, if some value D^*_s increases, we make a number of requests in point B_s for the associated matching problem: we make the same number of requests that the value D^*_s has been increased with.

We also associate an auxiliary matching algorithm (AMA) on this structure as follows. Suppose D^*_s increases at some time, causing jumps. These jumps are corresponding to requests of the associated matching problem; AMA satisfies these requests by the servers that are corresponding to those involved in these jumps.

Let $\hat{D}_s(p)$ denote the number of servers in block B_s just after phase p. If p is not the last phase, let $\varrho^{(p)+}$ denote the request sequence we get by adding the first request of phase $p+1$ to $\varrho^{(p)}$. Now we have

$$D_s^*(\varrho^{(p)}) \le \hat{D}_s(p) \le D_s^*(\varrho^{(p)+}) \tag{2}$$

and in all block but at most one we have equalities there (this is the block that causes termination of the pth phase).

Denote

$$m_p := \sum_{s=1}^{t} \max\{0, \hat{D}_s(p) - \hat{D}_s(p-1)\}. \tag{3}$$

Since the auxiliary metric space is uniform, the optimal cost is Δm_p.

Lemma 6 (Csaba, Pluhár, [8])**.** *The expected cost of AMA is at most* $\log k \cdot \Delta m_p$.

Lemma 7. *The expected cost of Algorithm X in the pth phase is at most*

$$f(k)\left(\sum_{s=1}^{t} \mathrm{opt}_s(D_s(\varrho^{(p)+}), \varrho^{(p)}) + \Delta\left(\sum_{s=1}^{t} D_s(\varrho^{(p)+}) - k\right)\right) +$$
$$+\Delta m_p(f(k)\log k + f(k) + \log k) + \Delta\frac{f(k)}{\log k}.$$

Proof. Consider the pth phase of a run of Algorithm X on the request sequence ϱ and let τ denote the associated run of AMA.

Let B_s be a block in which some request arrives during this phase. For the sake of convenience we omit the index s of the block: let $\mathrm{opt}(\ell, \varrho') := \mathrm{opt}_s(\ell, \varrho')$, $\hat{D}_{p-1} := \hat{D}_s(p-1)$ and $\hat{D}_p := \hat{D}_s(p)$.

Denote $\varrho^{(p)}$ the restiction of ϱ to B_s. While the block is unmarked, only jump-outs can happen from this block; let these jump-outs happen just before the r_1th, $\ldots, r_{d^-}$th request of $\varrho^{(p)}$, respectively. After the block has been marked, only jump-ins can happen; let these happen when the r_{d^-+1}th,…, $r_{d^-+d^+}$th request arrives, respectively (for any given request there can be more preceeding jumps). Denote $\sigma_i = \varrho_{r_i} \ldots \varrho_{r_{i+1}-1}$ (where ϱ_{r_0} is the first and $\varrho_{r_{d^-+d^++1}-1}$ is the last request of the phase in B_s), and let $k_{s,i} := k_i$ be the number of servers in B_s during σ_i. Observe that the demand at the r_ith request is exactly k_i. Finally, let ℓ_i denote the demand occuring at the $r_i - 1$th request if this request falls into the pth phase; otherwise let $\ell_i = 0$.

A jump-in to the block satisfies the last request, hence there is no server movement inside the block during a jump. The expected cost of non-jump movements in this block (this is called the *inner cost*) is, applying (1), at most

$$\begin{aligned}\sum_{i=0}^{d^-+d^+} \mathrm{E}[A_s(k_i, \sigma_i) | \, \tau] &\le \sum_{i=0}^{d^-+d^+} \left(f(k_i)\mathrm{opt}(k_i, \sigma_i) + \frac{k_i \cdot f(k_i)}{\log k_i}\delta\right) \\ &\le f(k) \sum_{i=0}^{d^-+d^+} \mathrm{opt}(k_i, \sigma_i) + \delta \sum_{i=0}^{d^-+d^+} \frac{k_i \cdot f(k_i)}{\log k_i}. \end{aligned} \tag{4}$$

We bound the right side of (4) piecewise. Summing up till the jump-out just before the last:

$$\sum_{i=0}^{d^- -1} \mathrm{opt}(k_i, \sigma_i) \le \sum_{i=0}^{d^- -1} \mathrm{opt}(k_{d^-}, \sigma_i) \le \mathrm{opt}(k_{d^-}, \varrho^{(p)}_{\le r_{d^-}}). \tag{5}$$

From the last jump-out till the last jump-in:

$$
\begin{align*}
\sum_{i=d^-}^{d-1} \mathrm{opt}(k_i,\sigma_i) &\le \sum_{i=d^-}^{d-1} \mathrm{opt}(\ell_{i+1},\sigma_i) \\
&= \sum_{i=d^-}^{d-1} \big(\mathrm{opt}(\ell_{i+1},\sigma_i) + \mathrm{opt}(\ell_{i+1},\varrho^{(p)}_{\le r_i}) - \mathrm{opt}(\ell_{i+1},\varrho^{(p)}_{\le r_i})\big) \\
&\le \sum_{i=d^-}^{d-1} \big(\mathrm{opt}(\ell_{i+1},\varrho^{(p)}_{<r_{i+1}}) - \mathrm{opt}(\ell_{i+1},\varrho^{(p)}_{\le r_i})\big) \tag{6} \\
&\le \sum_{i=d^-}^{d-1} \big(\mathrm{opt}(k_{i+1},\varrho^{(p)}_{<r_{i+1}}) + (k_{i+1}-\ell_{i+1})\Delta - \\
&\qquad\qquad - \mathrm{opt}(k_i,\varrho^{(p)}_{\le r_i}) - (k_i-\ell_{i+1})\Delta\big) \\
&\le \sum_{i=d^-}^{d-1} \big(\mathrm{opt}(k_{i+1},\varrho^{(p)}_{\le r_{i+1}}) - \mathrm{opt}(k_i,\varrho^{(p)}_{\le r_i}) + (k_{i+1}-k_i)\Delta\big) \\
&= \mathrm{opt}(k_d,\varrho^{(p)}_{\le r_d}) - \mathrm{opt}(k_{d^-},\varrho^{(p)}_{\le r_{d^-}}) + (k_d-k_{d^-})\Delta, \tag{7}
\end{align*}
$$

where $d = d^- + d^+$. Inequality (6) comes from Definition 4, since the demand of $\varrho_{<r_{i+1}}$ is ℓ_{i+1} and the demand of $\varrho_{\le r_i}$ is k_i.

Since $k_d \ge D(\varrho^{(p)})$, analogously we get

$$
\begin{align*}
\mathrm{opt}(k_d,\sigma_d) &\le \mathrm{opt}(D(\varrho^{(p)}),\sigma_d) \le \\
&\le \mathrm{opt}(D(\varrho^{(p)}),\varrho^{(p)}) - \mathrm{opt}(D(\varrho^{(p)}),\varrho^{(p)}_{\le r_d}) \le \\
&\le \mathrm{opt}(D(\varrho^{(p)+}),\varrho^{(p)}) + (D(\varrho^{(p)+}) - D(\varrho^{(p)}))\Delta - \\
&\quad - \mathrm{opt}(k_d,\varrho^{(p)}_{\le r_d}) - (k_d - D(\varrho^{(p)}))\Delta \\
&= \mathrm{opt}(D(\varrho^{(p)+}),\varrho^{(p)}) - \mathrm{opt}(k_d,\varrho^{(p)}_{\le r_d}) + (D(\varrho^{(p)+}) - k_d)\Delta. \tag{8}
\end{align*}
$$

Observe that if the request causing termination of phase p, that is the first request of phase $p+1$, then $D(\varrho^{(p)}) = D(\varrho^{(p)+})$ holds.

Summing up the right hand sides of (5), (7) and (8) we get

$$
\sum_{i=0}^{d^-+d^+} \mathrm{opt}(k_i,\sigma_i) \le \mathrm{opt}(D(\varrho^{(p)+}),\varrho^{(p)}) + (D(\varrho^{(p)+}) - k_{d^-})\Delta, \tag{9}
$$

and substituting this to the right hand side of (4) we get that the expected inner cost in B_s is at most

$$
f(k)\big(\mathrm{opt}(D(\varrho^{(p)+}),\varrho^{(p)}) + (D(\varrho^{(p)+}) - k_{d^-})\Delta\big) + \sum_{i=0}^{d^-+d^+} \frac{k_i\cdot f(k_i)}{\log k_i}\delta. \tag{10}
$$

On the other hand,

$$
D(\varrho^{(p)+}) - k_{d^-} = (D(\varrho^{(p)+}) - \hat{D}_p) + (\hat{D}_p - k_{d^-}), \tag{11}
$$

where we know that $D(\varrho^{(p)+}) = D_s(\varrho^{(p)+})$, and $(\hat{D}_p - k_{d^-})$ is the number of jump-ins into this block.

We can bound the sum of the expressions of the form $\frac{k_i \cdot f(k_i)}{\log k_i}\delta$ as follows:

$$\sum_{s=1}^{t}\sum_{i=0}^{d_s^-+d_s^+}\frac{f(k)}{\log k}k_{s,i}\delta \le \sum_{s=1}^{t}\sum_{i=0}^{d_s^+}\frac{f(k)}{\log k}k\delta \le (|\tau|+1)\frac{f(k)}{\log k}k\delta, \tag{12}$$

where $|\tau|$ is the number of the jumps during the phase. This comes from the fact that we have only k servers, hence the sum of k_s (the number of servers in the target block of the jump) and $k_{s'}$ (where the given server jumps from) is still at most k. We also remark that

$$|\tau| \le k. \tag{13}$$

Now we bound the cost of the jumps. Let T be the set of the potential jumps, and the length of the jump sequence η. Applying Proposition 6 we get that

$$\mathrm{E}[\eta\Delta] = \sum_{\tau\in T} P(\tau)|\tau|\Delta \le \log k \cdot \Delta m_p \tag{14}$$

Summing up the results (10), (11), (12) and (14) for all the blocks we get the following bound for the expected cost of Algorithm X:

$$\begin{aligned}
& \sum_{s=1}^{t}\sum_{i=0}^{d_s^-+d_s^+} \mathrm{E}[A_s(k_i,\sigma_i)+\eta\Delta] && (15)\\
= & \sum_{\tau\in T}\left(P(\tau)\sum_{s=1}^{t}\sum_{i=0}^{d_s^-+d_s^+}\mathrm{E}[A_s(k_i,\sigma_i)|\,\tau]\right)+\Delta\mathrm{E}[\eta] && (16)\\
\le & \sum_{\tau\in T}P(\tau)f(k)\sum_{s=1}^{t}\mathrm{opt}(D_s(\varrho^{(p)+}),\varrho^{(p)+})+ && (17)\\
& +\sum_{\tau\in T}P(\tau)f(k)\sum_{s=1}^{t}(D(\varrho^{(p)+})-\hat{D}_p)\Delta+ && (18)\\
& +\sum_{\tau\in T}P(\tau)f(k)|\tau|\Delta+\sum_{\tau\in T}P(\tau)(|\tau|+1)\frac{f(k)}{\log k}k\delta+\Delta\mathrm{E}[\eta] && (19)\\
\le & f(k)\left(\sum_{s=1}^{t}\mathrm{opt}_s(D_s(\varrho^{(p)+}),\varrho^{(p)})+\Delta\sum_{s=1}^{t}D_s(\varrho^{(p)+})-k\Delta+\log k\cdot\Delta m_p\right) \\
& +\Delta\left(\frac{f(k)}{\log k}m_p\log k+\frac{f(k)}{\log k}+m_p\log k\right),
\end{aligned}$$

if we apply $\sum_{\tau\in T}P(\tau)=1$ in (16) and (17), and $k\delta\le\Delta$ in (18). □

3.3 Analyzing the optimal cost

Consider an optimal solution of the k-server problem. Let $C_s^*(\varrho)$ be the maximal number of servers in B_s of this optimal solution during phase p and $C_s(\varrho)$ be the number of servers in B_s of the optimal solution at the end p. If we duplicate the first request of every phase but the very first one, and consider $\varrho^{(p)+}$ instead of $\varrho^{(p)}$, the optimal cost does not change.

Denote

$$ma_p^* := \sum_{s=1}^{t}(C_s^*(\varrho^{(p)+})-C_s(\varrho^{(p-1)+})).$$

From the definitions above it is clear that the optimal cost is at least

$$\mathrm{opt}(\varrho^{(p)+}) \geq \sum_{s=1}^{t} \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta \cdot ma_p^*. \tag{20}$$

Lemma 8.

$$\mathrm{opt}(k, \varrho^{(p)+}) \geq \sum_{s=1}^{t} \mathrm{opt}(D_s(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta \left(\sum_{s=1}^{t} D_s(\varrho^{(p)+}) - k \right).$$

Proof. From Definition 4 we have

$$\begin{aligned} &\sum_{s=1}^{t} \left(\mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \right) \geq \\ &\sum_{s=1}^{t} \left(\mathrm{opt}(D_s(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(D_s(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \right) \end{aligned}$$

Since $\sum_{s=1}^{t} C_s(\varrho^{(p-1)+}) = k$, the statement follows by (20). ☐

Lemma 9. *The optimal cost is at least*

$$\frac{1}{6} \Delta \sum_{p>1} m_p.$$

Proof. Let $\varrho^{(p)*}$ be the subsequence of $\varrho^{(p)}$ which we get by omitting each request that arrives to a block B_s after the demand of that given block reaches $D_s^*(\varrho^{(p)+})$ (note that $\varrho^{(p)*}$ is not neccessarily a prefix of $\varrho^{(p)}$). Now we have two cases: first, if $D_s^*(\varrho^{(p)+}) > C_s(\varrho^{(p-1)+})$ holds, then by Definition 4

$$\begin{aligned} & \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \\ \geq\ & \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)*}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \\ \geq\ & \left(\mathrm{opt}_s(D_s^*(\varrho^{(p)+}), \varrho^{(p)*}) + \Delta \left(0, D_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})\right) \right) \\ \geq\ & \Delta \max\{0, \hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))\}. \end{aligned}$$

Otherwise it holds that

$$\max\{0, \hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))\} = 0,$$

and also obviously

$$\mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \geq 0.$$

From both cases we get

$$\begin{aligned} \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \quad &\geq \\ \Delta \max\{0, \hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))\}. & \end{aligned} \tag{21}$$

Now since $\sum\limits_{s=1}^{t} \hat{D}_s(\varrho^{(p)}) = \sum\limits_{s=1}^{t} C_s(\varrho^{(p-1)+}) = k$, it also holds that

$$\begin{aligned} \sum_{s=1}^{t} \Delta \max\{0, \hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))\} \quad &= \\ \sum_{s=1}^{t} \frac{1}{2} \Delta |\hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))|. & \end{aligned} \tag{22}$$

Summing the cost of the jumps that the optimal solution performs we get

$$\Delta \sum_{s=1}^{t} |C_s(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})| \leq 2 \cdot \mathrm{opt}(k, \varrho^{(p)+}). \tag{23}$$

Note that the factor of 2 comes from the fact that each jump appears twice on the left hand side. Now summing up (21), (22) and (23) we have

$$\begin{aligned} & 4 \cdot \mathrm{opt}(k, \varrho^{(p)+}) \\ \geq \; & \Delta \sum_{s=1}^{t} \big(|\hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))| + |C_s(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})| \big) \\ \geq \; & \Delta \sum_{s=1}^{t} |\hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p)+}))|. \end{aligned} \tag{24}$$

Now summing (21) relativized to phase p and (24) relativized to phase $p-1$ we get that

$$\begin{aligned} & 2 \cdot \mathrm{opt}(k, \varrho^{(p)+}) + 4 \cdot \mathrm{opt}(k, \varrho^{(p-1)+}) \\ \geq \; & \Delta \sum_{s=1}^{t} \big(|\hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))| + |\hat{D}_s(\varrho^{(p-1)}) - C_s(\varrho^{(p-1)+}))| \\ \geq \; & \Delta \sum_{s=1}^{t} |\hat{D}_s(\varrho^{(p)}) - \hat{D}_s(\varrho^{(p-1)}))| = \Delta m_p, \end{aligned} \tag{25}$$

and the statement follows. □

Lemma 10. *If the pth phase is not the last one, then* $\mathrm{opt}(k, \varrho^{(p)+}) \geq \Delta$.

Proof. Similarly to the proof of Lemma 9,

$$\begin{aligned} & \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)+}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \\ \geq \; & \mathrm{opt}_s(C_s^*(\varrho^{(p)+}), \varrho^{(p)*}) + \Delta(C_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \\ \geq \; & (\mathrm{opt}_s(D_s^*(\varrho^{(p)+}), \varrho^{(p)*}) + \Delta \left(0, D_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+}))\right) \\ \geq \; & \Delta \max\{0, \hat{D}_s(\varrho^{(p)}) - C_s(\varrho^{(p-1)+}))\}, \end{aligned}$$

so we get that the optimal cost is at least

$$\sum_{s=1}^{t} \big(\mathrm{opt}_s(D_s^*(\varrho^{(p)+}), \varrho^{(p)*}) + \Delta(D_s^*(\varrho^{(p)+}) - C_s(\varrho^{(p-1)+})) \big). \tag{26}$$

Now since $\sum_{s=1}^{t} C_s(\varrho^{(p-1)+}) = k$, moreover if the pth phase is not the last one, then $\sum_{s=1}^{t} D_s^*(\varrho^{(p)+}) > k$ also holds; applying these we get the statement. □

3.4 Proof of Theorem 5

Now we are able to prove the theorem about competitiveness of Algorithm X.

Proof. [Theorem 5] If we apply (13) to the first phase and write k instead of $m_p \log k$, then Lemmas 7, 8, 9 and 10 give that

$$\begin{aligned} \mathrm{E}(\mathrm{cost}(KA(\varrho))) \; \leq \; & f(k)\mathrm{opt}(k, \varrho) + \\ & 6 \cdot \mathrm{opt}(k, \varrho)\big(f(k) \log k + f(k) + \log k\big) + \\ & \Delta \frac{f(k)}{\log k} \\ \leq \; & f'(k) \cdot \mathrm{opt}(k, \varrho) + \frac{f'(k) \cdot k\Delta}{\log k}, \end{aligned}$$

where $f'(k) = f(k)(6\log k + 8)$. □

At this point we make a few remarks:

1. If the starting configuration of the online problem coincides with the starting configuration of an optimal solution, then it holds that $\mathrm{E}(KA(\varrho)) \leq c \cdot f(k) \cdot \log k \cdot \mathrm{opt}(k, \varrho)$.
2. It is a bit more natural to require $\Delta \geq \delta M$ to hold, where M is the size of the greatest block. If additionally $M < k$ holds, we get a better competitive ratio.

3.5 Corollaries

Starting from the MARKING algorithm [10] and iterating Theorem 5 we get the following result:

Corollary 11. *There exists a $(c_1 \log k)^h$-competitve randomized online algorithm on any μ-HST of height h (here $\mu \geq k$), where c_1 is a constant. Consequently, when $h < \frac{\log k}{\log c_1 + \log\log k}$, this algorithm is $o(k)$-competitive.*

Considering only metric spaces having at most $t < k$ blocks, it is enough to require $\mu > t$. Substituting this to Lemma 7 we get a competitive ratio of $c \cdot \log b \cdot f(k)$, what (applying the above corollary) gives us the following result:

Corollary 12. *Suppose the metric space is a μ-HST having only degrees of at most b. Then there exists a c_2^h-competive randomized online algorithm, where h is the depth of the tree and $c_2 = c_2(b)$ is some constant.*

4 Further questions

In the field of online optimization the concept of buying extra resources is also investigated [9]. The quantity $\min_\ell\{\mathrm{opt}_s(\ell, \varrho) + \ell\Delta\}$ can be seen as the optimal cost of a model where one has to buy the servers, for a cost of Δ each. This problem on uniform spaces was studied in [9]. In this case $D_s(\varrho)$ is the number of servers bought in an optimal solution. Now considering to sequence ϱ_i, the behaviour of the associated sequence $D_s(\varrho_i)$ is unclear at the moment. It is an interesting question whether the above sequence is monotonically increasing, or does it hold that $|D_s(\varrho)_i - D_s(\varrho)_{i+1}| \leq 1$ for each i. The presented proofs would substantially simplify in both cases.

Another interesting question is that whether the $\log k$ factor in the competitive ratio per level of the HST is unavoidable, or an overall competitive ratio of $\Theta(\log k)$ holds for any HST.

Acknowledgement. The author wish to thank Béla Csaba for his guidance and suggestions, furthermore Péter Hajnal, Csanád Imreh and András Pluhár for their valuable remarks.

References

[1] Y. Bartal, Probabilistic approximation of metric space and its algorithmic application. *37th Annual Symposium on Foundations of Computer Sciences*, pp. 183–193, October 1996

[2] Y. Bartal, On approximating arbitrary metrics by tree metrics. *Proceedings of the 30th ACM Symposium on Theory of Computing*, pp. 161–168, 1998

[3] Y. Bartal, B. Bollobás, m. Mendel, A Ramsey-type theorem for metric spaces and its application for metrical task system and related problems. *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, Las Vegas, Nevada, 2001

[4] Y. Bartal, M. Mendel, Randomized k-server algorithms for growth-rate bounded graphs. *Journal of Algorithms*, **55(2)**:192–202, 2005

[5] S. Ben-David, A. Borodin, R. Karp, G. Tardos, A. Wigderson, On the power of randomization in on-line algorithms. *Algorithmica*, **11**:2–14, 1994

[6] A. Blum, H. Karloff, Y. Rabani, M. Saks, A decomposition theorem and bounds for randomized server problems. *Proceedings of the 33rd Symposium on Foundations of Computer Science*, pp. 197–207, October 1992

[7] B. Csaba, S. Lodha, A randomized on-line algorithm for the k-server problem on a line. *Random Structures and Algorithms*, **29(1)**:82–104, 2006

[8] B. Csaba, A. Pluhár, A randomized algorithm for on-line weighted bipartite matching problem. Will appear in *Journal of Scheduling*

[9] J. Csirik, Cs. Imreh, J. Noga, S. S. Seiden, G. Woeginger, Buying a constant competitive ratio for paging, *Proceedings of ESA 2001, LNCS 2161*, 98–108, 2001

[10] A. Fiat, R. M. Karp, M. Luby, L. McGeoch, D. Sleator, N. E. Young, Competitive paging algorithms. *Journal of Algorithms*, **12(4)**:685–699, 1994

[11] J. Fakcharoenphol, S. Rao, K. Talwar, A tight bound on approximating arbitrary metrics by tree metrics, *J. Comput. System Sci.* **69(3)**, 485–497, 2004

[12] E. Koutsoupias, C. Papadimitriou, On the k-server conjecture. *Journal of the ACM*, **42(5)**:971–983, 1995

[13] L. McGeoch, D. Sleator, A strongly competitive randomized paging algorithm. *Algorithmica*, **6(6)**:816–825, 1991

[14] M. Manasse, L. McGeoch, D. Sleator, Competitive algorithms for server problems. *Journal of Algorithms*, **11(2)**:208–230, 1990

[15] S. S. Seiden, A general decomposition theorem for the k-server problem. *Information and Computation*, **174(2)**:193–202, 2002